CSE 151B/251B: Deep Learning

Winter 2024

T-TH: 2-3:20 Warren Lectures Halls, room 2001

Section: 4-5PM, same location

Piazza page for the class is here

Canvas page for the class is here

Professor: Gary Cottrell

- gary@ucsd.edu (a very unreliable way to reach me!)
- Cell: 619-823-3033 (for really urgent matters: please use sparingly!)
- The best way to reach me is to email me and then text me to tell me to look for your email.
- Office Hours: I will hold office hours over zoom, for safety reasons (COVID is expected to surge again). (time to be determined) and Wednesdays at noon-1PM, or by appointment. Just go to my schedule for the zoom link - I won't be recording these. For appointments, go to my schedule and see when I am free for appointments. You have to use the week view to see the length of appointments. Pick a time, and email me a couple of suggestions. If you are able, sending me a google calendar invite is a good idea. **N.B.: I generally can't help with PyTorch or NumPy questions - only conceptual ones. I can look at your code and I can find bugs occasionally. But I don't actually program anymore!**

TAs: (click on a name to email)

For TA office hours, please see the class calendar, here.

Eric Yu (Head TA) **Shreya Sumbetla** **Ashrya Agrawal** **Andrew Ghafari**

Jay Jhaveri **Aishwarya Manjunath** **Naigam Shah**

Tutors: (should be used by 151B students): Tutors will be mostly used to help grading, so please don't contact them frequently.

Yuelei (Tina) Li **Yunfan Long** **Raj Nawal** **Buwei Wu**

COVID: It ain't over, and this s&^t is dangerous. You can recover and relapse with worse symptoms; you (very rarely so far) can become psychotic; you can end up in a "brain fog." Be careful, and stay safe!!! Call your loved ones.

Note: There will be no final; instead, there will be two midterms. Instead of a final, we will have final project presentations on Thursday, March 21st from 3-6PM (finals week).

Brief Class Description

Neural networks have come back into fashion since 2012 - especially for computer vision - when a deep network won the ImageNet Large Scale Visual Recognition Challenge (ILSVRC). They are now also used for language translation (google translate), image generation, game playing, and speech recognition. Most recently, they have been used to create chatbots (ChatGPT) and to create images from text descriptions (DALL-E 2 and friends). Now these capabilities have been combined in single large models.

In this course, we cover the **fundamentals** of neural networks: perceptrons, logistic regression, multilayer networks and back-propagation, deep convolutional neural networks, recurrent networks, transformers, and reinforcement learning. The course will involve four programming assignments roughly every two weeks, two midterms, and a final project. This course, while quite time consuming and difficult, is **not** an advanced deep learning course. We barely have time to cover the fundamentals! I do try to cover "what happened last week" in the last week of class.

We don't use books per se, but here are two resources: The first is Chris Bishop's 1995 book, Neural networks for Pattern Recognition, which can be found on the resources page as a pdf. This text contains a solid introduction to pattern recognition beyond just neural nets, especially the underlying statistical foundation. The text covers traditional pattern recognition, probability density estimation, single and multiple layer networks. Since the book is from 1995, well before deep networks became popular, the other resource is a recent interactive text called Dive into Deep Learning that I have not read (all of) yet, but it appears reasonably well written and it is interactive! Also, there will be some readings posted on Piazza under Resources.

Canvas

The syllabus, office hours, etc. will be there, but they will also be on Piazza. If there is conflicting information, Piazza is the standard, because we just copy stuff from Piazza to canvas.

No Zoom! But podcast!

All lectures are in person, but they will be podcast. Recordings will be available at podcast.ucsd.edu. This class is not listed there yet, but will be.

Clickers

We will be using clickers in order to check for understanding and for peer instruction purposes. Research has shown that discussing answers to questions in class with classmates leads to better understanding. Clicker use will be graded on participation only, and will only count for extra credit. We will use this to bump you up a grade (e.g., from A- to A) if you participate 75% of the time.

Piazza

We will be using Piazza as our class web board. You'll be able to ask the professor, the TAs and your classmates questions on Piazza. However, ***don't post your code, either as part of a question or part of an answer. Also, don't post your code in a public github repository!*** Thanks! Programming assignments and announcements will be posted there.

All of last year's lectures are posted there, along with supplementary reading. This year's lectures will also be posted there, usually **after** class, because I am always trying to optimize my slides. However, new slides typically will be very similar to old slides, and you can use those to follow along in lecture, or peruse before lecture.

Be sure to check your email or Piazza regularly.

ChatGPT, CoPilot and Friends

You **can** use generative AI models in your assignments to write code or your reports, but if you do, you are **required** to:

- Clearly mark where it was used in your report and/or code
- Add an appendix to your report that shows:
  - What prompt you used or the comments you used in Copilot autocomplete
  - The output of the AI and how you changed it for your problem/report. (If you didn't change it, just give the output of the AI and write "No Change"

Required coursework

The required work for this course will consist of two midterms (On Saturday, February 10th, and Saturday, March 9th, 10AM-Noon, in WLH 2001, as it says in the schedule of classes), and *five* programming assignments, at a brisk pace. Many of the programming assignments will include written homework. ***Hence you will need to devote significant effort to this course.*** Your programming assignments will count for 80% of your grade; the two midterms 10% each. There is usually an individual "homework" part that you should work on yourself. **Homeworks and PAs may include extra work for the grad students.**

The programming assignments will be due as follows:

PA 1: (paired assignment, 15% of grade) softmax regression and backprop

Handed out: Thursday, January 11th; Due: Thursday, January 25th.

PA 2: (groups of 3-4 people assignment, 15% of grade) Convolutional Networks

Handed out: Thursday, January 25th; Due: Thursday, February 8th.

PA 3: (groups of 4-5 people, 15% of grade): LSTM (Recurrent nets)

Handed out: Thursday, February 8th; Due: Tuesday, February 20th. (NOTE: Only 12 days for this one!)

Midterm, Saturday February 10th (review session the night before). (More details below)

PA 4: (groups of 4-6 people, 15% of grade): Transformers

Handed out: Tuesday, February 20st; Due: Tuesday, March 5th.

PA 5: (groups of 5+ people 20% of grade): *Final project of your choice*. See "Final Project Guidelines" on the resources page.

Handed out: Tuesday, March 5th.

Project proposal due **Thursday, March 7th**; (this will just be a few paragraphs so we can check that what you propose is reasonable);

Midterm, Saturday March 9th (review session the night before). (More details below)

Progress report due **Tuesday, March 12th.** (at this point, you should have your data, and be training networks).

Draft final report due **Sunday, March 17th**. (This should be nearly complete text of your report, perhaps with final result tables missing some data)

Project presentations: Thursday, March 21st, 3-6PM.

Final report due the next day, Friday March 22nd.

For all group work, whether the programming assignment or the final project, **every member of the group must include a (short) paragraph on what their contribution to the project was in the report. NO EXCEPTIONS.** Students who don't contribute may get a lower grade than other team members. ***Everyone should contribute to coding at least some part of the assignment***.

Course schedule with Readings

(D2L refers to Dive into Deep Learning; B95 refers to Bishop, 1995. For example, B95: 3.1-3.2 refers to Chapter 3, sections 1 and 2 in Bishop).

Week 0: Brush up on your math (optional): D2L: 2.3, 2.4 and 2.6 reminds you about linear algebra, calculus, and probability. **(**D2L:19 "Mathematics for Deep Learning" goes into more depth than you'll probably need.) See also B95 Ch's 1 and 2.

Week 1: Single-layer networks: l**ogistic regression** (D2L: 11.3.1-11.3.2; B95: 3.1.3), **softmax regression** (D2L: 3.4-3.7; B95: 3.1-3.2, 6.1, 6.7, 6.9), and **backprop** (B95: 4.1-4.3, 4.8; D2L: 5.1,5.3, 12.3 – although the way D2L talks about it is so different from the way I do, it may not be helpful. Here it might be better to stick with Bishop, but if you skip a lot of the math, it gives you the basic idea.);

Week 2: Representations (see Chapter8_with_notes_from_GWC, hinton86, Sejnowskispeech1987 under Readings - **if not there, please remind me to post them!**). **Tricks of the trade** (D2L: 12.4.1-12.4.2, 12.6.1, see lecun_98_efficient_backprop.pdf in Readings, Batch_Normalization.pdf in Readings); I**mproving generalization** (D2L: 5.5-5.6).

Week 3: Convolutional Networks (D2L: Ch 7; also see Stanford convnets course CS231n, Chapter8_with_notes_from_GWC pp. 348-352 in Readings, lecun89e in Readings).

Week 4: More Convnets (D2L: Ch 8 – I certainly won't cover everything in this chapter, but it's a good reference.) **Beginning of Recurrent Nets** (D2L: 9.1,9.4,9.7 Chapter8_with_notes_from_GWC pp. 354-361, papers by Jordan, St. John, Elman under Readings); LSTM networks (D2L: Ch 10 – try to skip over the math! ;-))

Review session Friday, February 9th, 7-9PM (will be recorded)

Midterm Saturday, February 10th, 10:00PM-12PM, (ROOM TBA)

Week 5: Recurrent networks (D2L: 9.1,9.4,9.7 Chapter8_with_notes_from_GWC pp. 354-361, papers by Jordan, St. John, Elman under Readings); LSTM networks (D2L: Ch 10 – try to skip over the math! ;-))

Week 6: Attention networks (D2L: 11.1, 11.3, 11.5) more applications (see DRAM and neural turing machine paper in readings (remind me to post if not there);

Week 7: Transformers (D2L: 11.7)

Week 8: Reinforcement Learning See lecture notes.

Week 9: More Reinforcement Learning; **GANs** (D2L:18.1)

Review session Friday, March 8th, 7-9PM (will be recorded)

Midterm Saturday, March 19th, 10:00AM-NOON, (Room TBA) Cumulative, but emphasis on second half of course.

Week 10: Selected recent papers; Ethics of AI

First assignment

The first reading assignment is to read D2L Ch 1, and B95 Ch 1 for general background (we will not cover this in any detail, but it is good for you to know). The first programming assignment will be posted on Piazza, and will include some (individual) written work.

Resources

We will endeavor to answer your questions on Piazza in a timely manner. We will hold office hours and/or sections almost every day of the week. There are many online deep learning courses that have good lectures available. If you are not getting answers to your questions on Piazza within 24 hours, please text the professor at 619-823-3033 and I will give the TAs holy hell.

- There are courses online where the giants in the field explain neural networks and deep learning. Geoff Hinton had a coursera neural network course; many of the lectures are available on Youtube if you search for "youtube hinton neural network lecture" – they are the ones that say "Lecture 1.1", etc. He also has a page of tutorials here.
- Andrew Ng has a coursera course on deep learning here. You can sign up for a 7-day free trial.
- Stanford's convolutional network course is here. The initial web page has a convolutional network running in your browser! The 2017 lectures are free and on youtube.
- Andrej Karpathy's lectures for that course from Winter 2016 are here.
- The neural network playground is a great tool to play around with a neural network and get insights into the kinds of features they learn. You can vary the number of hidden layers, the activation functions, the learning rates, etc.
- Our demo of face recognition is somewhat amusing, but requires matlab.
- Andrej Karpathy's blog is just wonderful. There are lots of other blogs also.
- Lilian Weng's blog has explanations of almost everything!
- Two minute papers are an excellent way to keep up, without *really* keeping up…

Meta-resource:

- Learning How to Learn, the most popular MOOC there is. Taught by our own Terry Sejnowski and not-our-own Barbara Oakley.

Grading policy

In graduate classes, I curve to a B+, unless doing so would lower your grade! My procedure is as follows: I take the median grade, add sufficient number of points to make it 88.5 (middle of a B+), add that number to all grades, and then round **up** the result to the nearest whole number. I use standard cutoffs, except at the top and the bottom, where I make the cutoff for an A+ 96 instead of 97, and at the bottom, I use Charles Elkan's dictum that "anyone who learns half of what I teach them deserves to pass." After deciding the grades this way, I use your participation points to raise your grade by a notch - e.g., if you have an A-, it will be raised to an A ***if you have 75% participation.***

Late Policy

Assignments can be turned in a day late with a 10% penalty. Two days late: 20%. Three days late: 50%. More than three days: 100%. These can be extended due to illness or having to tend to someone ill.

Academic Integrity

Working in small groups in programming assignments is required. However, on written homeworks (derivations and such), please follow the **Gilligan's Island rule**:

(Dymond, 1986): No notes can be made during a discussion, and you must watch one hour of Gilligan's Island or some equally insipid TV show before writing anything down. Suspected cheating has been reported to the Dean in the past, and will be again.

On midterms, you may bring one 8.5"x11" page of notes, one-sided. Of course, please do not read off of other students' tests, etc. We have had complaints of cheating in the past during tests with graduate students. It seems unbelievable, but it happens. We will be watching closely. Anyone who is caught cheating on a test will receive a 0 for that test. We will give out two versions of each test.

To detect instances of academic integrity violations in programming assignments we will use 3rd party software. We recommend you do not put your name in your code, unless you don't mind this potential exposure. Including your name and/or PID will disclose that information to the 3rd party.

Prerequisites

This course has no formal prerequisites, but some mathematical sophistication is required. You should know some probability, and know how to sample from a distribution. You should know linear algebra, and vector calculus (at least, partial derivatives). This course is intended for CSE, Cognitive Science and ECE graduate students. The use of numPy is strongly encouraged for the first two assignments. After those two, we will switch to using a deep learning platform – we will use PyTorch because it has good facilities for inspecting internal variables, which can aid in debugging.

Texts

Neural Networks for Pattern Recognition, by Chris Bishop, Oxford University Press, 1995.

Dive into Deep Learning by Aston Zhang, Zachary C. Lipton, Mu Li and Alexander J. Smola, 2020.

Supplementary reading:

Duda, Hart and Stork *Pattern Classification* (2nd Ed). Wiley

For the Physics inclined:
Introduction to the Theory of Neural Computation. Hertz, Krogh & Palmer, Addison Wesley.

For the Cognitively iclined:
Parallel Distributed Processing, Vols 1 & 2, edited by Rumelhart and McClelland (MIT Press, 1986) and *Explorations in PDP*, by McClelland & Rumelhart (MIT Press, 1987) (This is a good way to get a working introduction to neural nets). These are available as pdf's from Jay McClelland's website.

Diversity and Inclusion

We are committed to fostering a learning environment for this course that supports a diversity of thoughts, perspectives and experiences, and respects your identities (including race, ethnicity, heritage, gender, sex, class, sexuality, religion, ability, age, educational background, etc.). Our goal is to create a diverse and inclusive learning environment where all students feel comfortable and can thrive.

Our instructional staff will make a concerted effort to be welcoming and inclusive to the wide diversity of students in this course. If you have a particular gender pronoun you prefer, or some other way we can make you feel more included please let one of the course staff know, either in person, via email/discussion board, or even in a note under the door. Our learning about diverse perspectives and identities is an ongoing process, and we welcome your perspectives and input.

We also expect that you, as a student in this course, will honor and respect your classmates, abiding by the UCSD Principles of Community (https://ucsd.edu/about/principles.html). Please understand that others' backgrounds, perspectives and experiences may be different than your own, and help us to build an environment where everyone is respected and feels comfortable. In particular, be respectful in responding to other students on the Piazza page.

If you experience any sort of harassment or discrimination, please contact the instructor as soon as possible. If you prefer to speak with someone outside of the course, please contact the Office of Prevention of Harassment and Discrimination: https://ophd.ucsd.edu/.

Students with Disabilities

We aim to create an environment in which all students can succeed in this course. If you have a disability, please contact the Office for Students with Disability (OSD), which is located in University Center 202 behind Center Hall, to discuss appropriate accommodations right away. We will work to provide you with the accommodations you need, but you must first provide a current Authorization for Accommodation (AFA) letter issued by the OSD. You are required to present their AFA letters to Faculty (please make arrangements to contact me privately) and to the OSD Liaison in the department in advance so that accommodations may be arranged.

Basic Needs Resources

Are you eating properly? Do you have adequate access to nutritious food? Do you have stable housing? Are you homeless or couch surfing? If you or someone you know has food and/or housing insecurity, please note: http://basicneeds.ucsd.edu

The Triton Food Pantry (in the old Student Center), https://www.facebook.com/tritonfoodpantry/ is free and anonymous, and includes produce.

Financial aid resources, the possibility of emergency grant funding, and off-campus housing referral resources are available. CAPS and college deans can connect students to the above resources, as well as other community resources and support.